arXiv:1512.02013v1 [cs.CV] 7 Dec 2015

# SCALABLE DOMAIN ADAPTATION OF CONVOLUTIONAL NEURAL NETWORKS

*Adrian Popescu, Etienne Gadeski, Hervé Le Borgne*

CEA, LIST, Vision and Content Engineering Laboratory
France
Email: {adrian.popescu,etienne.gadeski,herve.le-borgne}@cea.fr

**ABSTRACT**

Convolutional neural networks (CNNs) tend to become a standard approach to solve a wide array of computer vision problems. Besides important theoretical and practical advances in their design, their success is built on the existence of manually labeled visual resources, such as ImageNet. The creation of such datasets is cumbersome and here we focus on alternatives to manual labeling. We hypothesize that new resources are of uttermost importance in domains which are not or weakly covered by ImageNet, such as tourism photographs. We first collect noisy Flickr images for tourist points of interest and apply automatic or weakly-supervised reranking techniques to reduce noise. Then, we learn domain adapted models with a standard CNN architecture and compare them to a generic model obtained from ImageNet. Experimental validation is conducted with publicly available datasets, including Oxford5k, INRIA Holidays and Div150Cred. Results show that low-cost domain adaptation improves results compared to the use of generic models but also compared to strong non-CNN baselines such as triangulation embedding.

***Index Terms***— convolutional neural networks, adaptation, visual resources, scalability

## 1. INTRODUCTION

Many computer vision tasks, including object classification and localization, as well as content based retrieval, are increasingly tackled with convolutional neural network (CNN) architectures. A lot of recent papers [1, 2] focus on improving CNN architectures and assume that large-scale manually labeled datasets, such as ImageNet [3], are readily available. However, while large, the coverage of ImageNet is insufficient in many domains. For instance, it illustrates only a very limited amount of named entities, including tourist points of interest or car brands and models. Such entities can nevertheless be of interest in domain-related applications and good quality visual resources which illustrate them are assumed to be necessary. Manual resource enrichment is tedious and not considered as scalable in terms of number of concepts and images. To overcome this issue, this paper focuses on domain adaptation and more precisely on automatic or weakly-supervised methods to create the visual resources necessary for domain specific CNN training.

Due to resource scarcity, domain transfer recently received a particular attention from the computer vision community. The authors of [4] have shown that activations of top layers of a CNN can be used effectively for image retrieval. More interestingly, they discovered that performance increases when the network is retrained with images similar to a specific domain (tourist points of interest). However, their method requires important manual intervention, since they verify 200 images for each tourist point of interest (POI) included in the training set. Equally important, they assume that the presence of noisy images reduces the quality of CNN models and create an unbalanced training set (i.e. some POIs are illustrated by 1000 images, while a majority of them are represented by only 100 images). In this paper, we focus on methods to build a CNN training set in a fully automated manner and show that, contrary to the hypothesis of [4], a CNN can actually be trained from scratch with these images. This is probably due to the intrinsic quality of the training dataset constituted with our method but also to its larger size.

The automatic collection of groups of visually coherent images was already addressed in literature. The general idea consists of collecting "noisy" images which are then reranked according to a learned model [5] or use clustering to determine visually compact groups [6]. Beyond the fact these works did not aim at determining groups of image for CNN learning, we propose a simpler method, based on a finer domain characterization. Hence, our approach has a better computational tractability, while preserving CNN model accuracy.

Experiments are carried out on three publicly available datasets which include images related to the tourism domain but have different visual properties. We test with Oxford5k [7], INRIA Holidays [8] and Div150Cred [9] and show we obtain better results than other CNN-based method and among the best reported in the literature on these sets.

## 2. RELATED WORK

CNNs have recently shown impressive performances in different image mining tasks [10, 11]. For instance, the use

of CNN ImageNet ILSVRC challenge reduced error rate to 0.15 in 2012, 0.11 in 2013 and 0.067 in 2014 [12] from 0.26 obtained with Fisher Vectors [13], the best performing pre-CNN features. A wide majority of CNN related work builds on the architecture proposed in [10] and proposes different improvements. For instance, the best performing entries of ILSVRC 2014 [14, 15] focus on model depth and give priority to different optimizations meant to reduce the computation time. In parallel, a lot of attention is paid to the way available data are exploited. Domain transfer is another active area. The authors of [16] exploit CNN mid-level features learned on ImageNet data and adapt two full-connected layers to successfully classify PascalVOC 2007 data. In a cruder approach, [17] test off-the-shelf `OverFeat` features [2] on a wide series of benchmarks. They show that CNN features constitute a competitive (when not state-of-the-art) approach compared to other features which are optimized for particular datasets. Data augmentation through simple image transformation is carefully studied in [1] and leads to consequent performance improvement on the PascalVOC 2007 dataset. The use of the spatial structure of the image is successfully introduced in [17] as a simple multiscale feature extraction and max-pooling scheme. A more elaborate scheme is proposed in [18], which exploits multiscale extraction, VLAD pooling and concatenation. Both [17] and [18] report important improvement on tourism related datasets, such as INRIA Holidays and/or Oxford5k. The work closest to ours is presented in [4], where authors also train a network dedicated to the tourism domain. However, their approach differs from ours in several important aspects: (1) the amount of supervision tested is significantly higher; (2) no image reranking is tested to reduce the influence of noise; (3) training is done on an highly unbalanced dataset, a modeling choice which is likely to downgrade results; (4) they initialize with the weight learned on ImageNet; (5) the data sources are different (Yandex in [4] vs. Flickr here). One common assumption of most CNN-related works cited above is to assume the pre-existence of a manually built visual resource good enough to train CNNs. In most cases, this resource is a subset of ImageNet, which is used to classify images of the same concepts [12] or to classify images from other datasets [17, 16, 1]. Departing from the assumption that training data are readily available, we focus on (semi)automatic visual resource creation. This question was already studied in non-CNN context [5, 6] but, to our knowledge, only seldom in a CNN context [4]. This research question has important implications because, if such resources can be built with minimal manual effort, image classification could be easily applied to a large number of domains, provided that enough data are available to describe associated concepts.

Since the early works on efficient image retrieval [7], many refinement were proposed such as query expansion [19], more efficient features [20] and aggregating local features efficiently [21]. Some of the best results were recently obtained with triangulation embedding [22], that are particularly efficient in their compressed version, which enables scalable image search at the cost of a limited loss of performance.

## 3. CNN DOMAIN ADAPTATION APPROACH

CNNs are majoritarily used for solving computer vision problems in which train and test concepts are the same [12]. However, if trained with large arrays of concepts, generic feature extractors can also be used to characterize other datasets whose concepts overlap the original ones to some extent [16, 17]. Transfer efficiency is reduced whenever the gap between train and test sets is too high and, in such cases, dedicated models should be trained. In this section, we describe the main steps of our domain adaptation approach which consists of the following main steps: (1) deriving relevant training concept for a given domain; (2) collecting and, in the case of noisy data, reranking training images; (3) training a CNN model with adapted data. During the entire process, we assume that a generic CNN models is readily available for feature extraction and use the `Caffe` reference model [23] which is a slightly modified version of the one proposed in [10].

### 3.1. Domain characterization

One central working hypothesis here is that, when performing domain transfer with CNNs, train and test sets should be reasonably well linked. This link can be implemented through the use of similar concepts during training and testing. For instance, an image of the *Empire State Building* is more likely to be well described by a CNN model obtained with a significant number of other skyscrapers among the training concepts compared to a training done with ImageNet, which contains a very low number of skyscraper related concepts. Domain adaptation is realized through the use an external knowledge base which contains concepts from the target domain. For instance, if the test set contains a lot of tourism photographs, as it is the case here, we extract a list of POIs from an existing gazetteer [24] and be used as conceptual support for adapted CNN training. This approach assumes that an appropriate textual knowledge base is readily available for the target dataset(s). It would also be possible to take a data driven approach, in which images from the test set are compared to a large catalogue of visual models in order to characterize the domain. Such a visual catalogue could be obtained, for instance, from Wikipedia but this approach falls outside of the immediate scope of this work.

### 3.2. Data collection and reranking

The availability of large visual resources is a key component in CNN training and usage [12]. As we mentioned, when

dealing with large test sets, a subset ImageNet is the preferred training resource. While this resource ensures large conceptual coverage, it doesn't cover domains equally well. Tourism is one such domain and we investigate ways to characterize it with a (semi)automatically built resource. Since full manual annotation of domain visual resources requires consequent efforts, we focus on low-cost alternatives. More precisely, we investigate the following strategies:

- POI-CNN-N - direct used of Flickr image sets.
- POI-CNN-A - automatic reranking based on the correlations between the images associated to each POI.
- POI-CNN-W - weakly supervised reranking in which candidate images are compared against a limited amount of manual annotations.

POI-CNN-N and POI-CNN-A do not require any manual intervention and are most interesting from a scalability perspective. POI-CNN-A is implemented following the conclusions of [5], where the authors show that reliable models can be built from limited and noisy sets of positive images. In practice, positive examples were chosen from an initial random sample of 60 images of each POI. POI-CNN-W represents a compromise between no manual annotation and full dataset annotation, as done during the creation of ImageNet [3]. POI-CNN-A and POI-CNN-W strategies are implemented using a linear SVM based reranking method [5], applied to a generic CNN description of image content. The choice of the reranking method is done as a compromise between performance and scalability. To obtain a good separation of positive and negative samples, we use a 1:100 ratio between positive and negative samples. A unique negative class is defined from 500,000 diversified images. We use the $\ell^2$-regularized $\ell^2$-loss SVM implemented in liblinear, which is solved in the primal using a trust region Newton method. We determined the penalty parameter on an independent dataset (Pascal VOC) and fixed it to $C = 1$ for all experiments. For POI-CNN-A, SVM reranking is performed by using a 10-fold split of the initial positive set.

### 3.3. CNN training

We train CNN models using the different reranking strategies presented in the preceding subsection in order to assess their influence on the quality of results. POI-CNN-N, POI-CNN-A and POI-CNN-W are trained with raw Flickr images, automatically reranked images and weakly supervised reranked images respectively. To ensure comparability, the same set of POIs is used in all three cases, with the top ranked 1000 images kept for each POI. As we mentioned, CNN architecture optimization falls outside of our immediate scope here. Throughout this paper, we exploit the network architecture and training procedure proposed in [23] to obtain the `Caffe` reference models.

### 3.4. Feature tuning

Previous works [7, 19, 17] have shown that visual instance retrieval benefits from different feature tuning methods. However, many of these improvements, including geometric verification or recursive query expansion [19], imply high computational complexity and make them difficulty usable for large scale retrieval. Since our main objective is to propose efficient and scalable features for the tourism domain, we adopt only tuning methods which keep feature computation and retrieval manageable:

- spatial search (S) - simple procedure described in [17] in order to cope with object location and scale variability. Overlapping patches are extracted at different scales, with the size of each patch being $\frac{2}{l+1}\times$(width, height), where $l$ is the patch level. At each scale we extract $l^2$ patches. At even levels we also extract a patch at the center. The similarity between two patches is computed by retaining the minimum $\ell^2$ distance of the query patch to all patches of reference images. The similarity between a query and a reference image is the average of minimum distances between patches. Similar to [17], we test different scales for query and reference images. This step has important effect on retrieval complexity since it involves $h_q \times h_r$ comparison of query and reference image patches ($h_q$ is the number of levels for query images and $h_r$ for reference images). More advanced object localization techniques exist [15, 2] but their computational cost is too high in our setting.
- feature augmentation (A) - pipeline including $\ell^2$ normalization, PCA-base compression, feature whitening and $\ell^2$ re-normalization introduced in [25] and exploited in [17].
- query expansion (QE) - both average query expansion method [19] and discriminative query expansion [20] were initially tested and, since we didn't find any significant difference, the first one was retained due to its lower complexity. This method exploit pseudo-relevance feedback to update query representation based on an initial list of top results. Since spatial verification is not straightforward with CNN features, it is not exploited here.

## 4. EVALUATION

We evaluate different aspects of our approach using challenging and publicly available datasets related to the tourism domain:

- *Oxford5k* (Ox5k) [7] - includes photos for 11 different buildings from the Oxford area, with the buildings occupying a significant part of the photo for relevant

items. All buildings have a stable and distinctive appearance, making them suitable for geometric confirmation of descriptions.

- *Holidays* (Holi) [8] - is a good representative of a holiday photos folder since it includes both landscapes and POIs. Photographs are more diversified compared to *Oxford5k* but the dataset is smaller.

- *Div150Cred* [9] was very recently introduced in the context of the MediaEval 2014 evaluation exercise. The dataset includes high diversity of POI types, with variable visual complexity, is represented. *Div150Cred* includes more complex objects (*i.e.* museums, squares, parks) than *Oxford5k*. Consequently, relevant images of these objects are much more diversified compared to those of buildings and the dataset is more challenging.

The usual evaluation protocols are followed for *Oxford5k* and for *Holidays*. Our approach is compared to a selection of competitive methods proposed in literature, with particular focus on mid-sized descriptors. *Div150Cred* was used to assess reranking methods for intra-POI retrieval. We repurpose this dataset for inter-POI retrieval and use the 665 Wikipedia images provided as examples along with the collection as queries. Naturally, a result is considered relevant if it belongs to the same POI as a query image. To compare our approach with existing work on this dataset, we compute results with bags of visual words with a very large vocabulary, a competitive non-CNN method and also with the `Caffe` reference model, an off-the-shelf CNN feature extractor. Similar to the other datasets, performance is reported using Mean Average Precision (mAP)[1].

All adapted CNN models are learned using the network configuration proposed in [23][2]. The learning process is stopped after 350,000 iterations since no significant improvement is obtained during the previous 50,000 iterations. All results are reported with features extracted from the last fully connected layer (`fc7`). Learning was done using GTX Titan Black GPUs and took approximately one week for each model. Spatial search is performed with $h_q = 3$ and $h_r = 4$ for Oxford5k and with $h_q = 2$ and $h_r = 2$ for Div150Cred.

### 4.1. Overall results

The results obtained with the existing high-performance models and with the proposed approach are summarized in Table 1. The main finding is that the adapted CNN features introduced here perform better than generic `Caffe` features for all three datasets, a result which indicates that domain adaptation is effective. Differences are particularly high for *Oxford5k* and *Div150Cred* datasets, with improvement of over 100% and respectively 200% over the generic `Caffe` features. Equally important, we obtain state-of-the art performance among CNN on both Oxford5k and Holidays. Compared to the best non-CNN methods, adapted features still lag behind fine-tuned high-dimensional features [20], [21] for *Oxford5k* and has comparable performance on *Holidays*. It is noteworthy that the results reported in [20], [21] are obtained with features which are less scalable than ours. The adapted CNN features compare favorably with the best configuration of the recently proposed triangulation embedding [22], for both *Oxford5k* and *Holidays*. This comparison is interesting because the two types of features are comparable from a scal-

**Table 1**. mAP results obtained on *Oxford5k*, *Holidays* and *Div150Cred* datasets. T-Emb is the best configuration obtained with triangulation embedding in [22]. *Over* and *Caffe* results are obtained `OverFeat` [2] and `Caffe` [23] with the pretrained ImageNet CNN models provided along with these tools. $\text{CNN}_\text{d}$ is the domain adapted CNN representation introduced in [4]. For each existing work, we report only the best performance obtained for each dataset. S, A, QE stand for spatial search, feature augmentation and query expansion respectively. POI-CNN-N, POI-CNN-A, POI-CNN-W are our adapted CNN models with no reranking, with automatic and with weakly supervised reranking. When used with subscripts, the features of our three models are compressed with PCA and the number indicates the dimension of the resulting vector. For spatial search features dimension varies because it is related to the number of patches we consider for one image. For inverted index approaches, the dimension is that of the codebook.

| Method | Dim. | S | A | QE | Ox5k | Holi | DIV |
|---|---|---|---|---|---|---|---|
| Inverted index approaches | | | | | | | |
| [20] | 1M | | | x | **92.9** | N/A | N/A |
| [21] | 65k | | | | 83.8 | **88** | N/A |
| T-Emb [22] | 8k | | | | 67.6 | 77.1 | N/A |
| BoVW | 1M | | | | 72.9 | 51.2 | 2.7 |
| Existing CNN methods | | | | | | | |
| *Over* [17] | 4k | | | | 32.2 | 64.4 | N/A |
| *Over* [17] | 4k | x | x | | 68 | 84.3 | N/A |
| MOP-CNN [18] | 12k | | | | N/A | 78.8 | N/A |
| $\text{CNN}_\text{d}$ [4] | 4k | | | | 54.5 | **79.3** | N/A |
| *Caffe* [23] | 4k | | | | 38.2 | 73 | 3.7 |
| *Caffe* | 4k | x | x | | 71.2 | 87.1 | 6.4 |
| Proposed CNN methods | | | | | | | |
| POI-CNN-N | 4k | | | | 65.7 | 77 | 10.3 |
| POI-CNN-A | 4k | | | | 64.6 | 76.5 | 9.8 |
| POI-CNN-W | 4k | | | | 67.1 | 76.3 | 9.8 |
| POI-CNN-$\text{N}_{512}$ | 512 | | | | 68.6 | 78 | 11.9 |
| POI-CNN-$\text{A}_{512}$ | 512 | | | | 66.4 | 77.9 | 10.9 |
| POI-CNN-$\text{W}_{512}$ | 512 | | | | **69.3** | 77.7 | 10.8 |
| POI-CNN-$\text{N}_{512}$ | 3-16k | x | x | | **76.9** | **87.5** | **12.5** |
| POI-CNN-$\text{A}_{512}$ | 3-16k | x | x | | 76.7 | 86.9 | 11.5 |
| POI-CNN-$\text{W}_{512}$ | 3-16k | x | x | | 76.7 | 86.6 | 11.8 |

[1] A complete evaluation kit will be provided in order to facilitate reproducibility.

[2] The lists of POIs and of images used for learning and the network configuration files will be provided in order to facilitate reproducibility.

**Table 2**. Effect of PCA-based dimensionality reduction. Results are presented for adapted CNN features obtained without reranking (POI-CNN-N).

| | Dimension | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **4096** | **1024** | **512** | **256** | **128** | **64** | **32** | **16** |
| Ox5k | 65.7 | 67.8 | 68.6 | 69.2 | 68.6 | 66.2 | 62.2 | 53.9 |
| Holi | 77 | 76.8 | 78 | 78.5 | 77.8 | 75.5 | 71.3 | 63.9 |

ability point of view. BoVW with a very large vocabulary is another strong baseline and a comparison to it shows that the best configurations of adapted CNN features outperform it on all three test datasets. The gain is particularly important for Div150Cred, where the mAP score goes from 2.71% to 12.5%. This result is important from an application perspective since the Div150Cred dataset is the closest proxy to real Web tourism corpora.

Surprisingly, the results obtained with POI-CNN-N, *i.e.* direct use of Flickr images for POIs, are comparable, when not better, to those obtained with automatic image reranking and weakly supervised reranking (POI-CNN-A and POI-CNN-W, respectively). A qualitative evaluation of reranking results indicates that the two proposed methods are effective in cleaning noise but also decrease the overall diversity of the images. An empirical explanation of the comparable results of the three methods is that, up to a certain level, training set noise is compensated by diversity. In other words, a diversified negative set is at least as important as a clean positive set. If confirmed beyond the tourism domain, this finding has important implications since CNN training could be done without cumbersome annotation of large volumes of data.

### 4.2. Feature compression effect

We apply a standard PCA-based compression of the adapted CNN features in order to improve their scalability and present the obtained performance in Table 2 for *Oxford5k* and *Holidays* datasets. Overall, the best results are obtained when retaining the first 256 components of the PCA vectors and this shows that compression is actually beneficial when applied to adapted CNN features. Below 256 dimensions, performances start to drop but interesting results are obtained even when retaining as few as 16 dimensions. For instance, on Oxford5k, a mAP of 53.9% is similar to [4] (54.5%) for POI-adapted CNNs with 4k dimensions. The behavior of the features proposed here with respect to compression shows that they provide a good balance between scalability and efficiency and could thus be exploited to mine Web-scale datasets.

### 4.3. Feature tuning effect

Confirming results presented in [17], spatial search and data augmentation bring further improvement compared to the direct use of full image features. Relative mAP improvements of over 10% are obtained for *Oxford5k* and *Holidays* and between 5% and 10% for *Div150Cred*. The lower improvement for the latter dataset is explained by the fact that the initial image size is smaller (500 pixels on the larger dimension) and only two levels were used for spatial search. The best previous retrieval results are obtained with generic features, combined with spatial search and feature augmentation [17] using `OverFeat`. We repeated these experiments with `Caffe` and results reach $71.2\%$, $87.1\%$ and $6.4\%$ on *Oxford5k*, *Holidays* and *Div150Cred*, respectively. While generic feature tuning effectiveness gives good results on the first two datasets, it clearly fails to match the performance of domain adaptation on the more complex *Div150Cred*.

In addition to spatial search and data augmentation presented in Table 1, query expansion has a beneficial effect for *Oxford5k*. We adopt an average query expansion method using the top 20 results from an initial list as positive samples for building an improved query. We therefore reach a mAP of 81.5%, 80.9% and 82.7% for unsupervised, automatic reranking and weakly supervised reranking models respectively. However, this form of query expansion didn't have a positive effect on *Holidays* and *Div150Cred*.

## 5. CONCLUSIONS

We proposed a CNN adaptation pipeline which learns features from an domain-adapted but noisy Web dataset and successfully transfers them for retrieval in the target domain. One important contribution is to show that learning is effective even in presence of noisy images. If confirmed in other domains, this finding has important implications since it indicates that manually created resources, such as ImageNet, could be successfully replaced with Web datasets for image retrieval. Naturally, the quality of the obtained representations depends on an appropriate choice of domain concepts and on the quality of the underlying image source.

A second contribution is to show that domain adaptation of CNN features is efficient. Compared to [4], we obtain better results while requiring no manual supervision and also by a more careful setting of the learning process. In addition, our results regarding the usefulness of feature tuning confirm those of [17].

The results reported here are very encouraging and we intend to pursue work in different inter-related directions. Experiments were performed with datasets from the tourist domain and one important future work axis is to confirm them on other domains. Another is to test more complex CNN architectures [15] and to feed the network with larger datasets (*i.e.* more concepts and more images per concepts). Finally, building on [18], we want to explore the alteration of CNN architectures in order to better account for spatial context.

## 6. REFERENCES


[1] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional nets," in *British Machine Vision Conference*, 2014.

[2] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun, "Overfeat: Integrated recognition, localization and detection using convolutional networks.," *arXiv:1312.6229 [cs.CV]*, 2013.

[3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition*, 2009.

[4] A. Babenko, A. Slesarev, A. Chigorin, and V. Lempitsky, "Neural codes for image retrieval," in *European Conference on Computer Vision*, 2014.

[5] F. Schroff, A. Criminisi, and A. Zisserman, "Harvesting image databases from the web," *Pattern Analysis and Machine Intelligence*, vol. 33, no. 4, pp. 754–766, 2011.

[6] D. Tsai, Y. Jing, Y. Liu, H.A. Rowley, S. Ioffe, and J. M. Rehg, "Large-scale image annotation using visual synset," in *International Conference on Computer Vision*, 2011.

[7] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman, "Object retrieval with large vocabularies and fast spatial matching," in *Computer Vision and Pattern Recognition*, 2007.

[8] H. Jégou, M. Douze, and C. Schmid, "Hamming embedding and weak geometric consistency for large scale image search," in *European Conference on Computer Vision*, 2008.

[9] B. Ionescu, A. Popescu, M. Lupu, and H. Muller, "Retrieving diverse social images at mediaeval 2014: Challenge, dataset and evaluation," in *Working Notes MediaEval*, 2014.

[10] A. Krizhevsky, I. Sutskever, and G.E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Adv. in Neural Information Processing Systems*, 2012.

[11] D. C. Ciresan, U. Meier, and J. Schmidhuber, "Multi-column deep neural networks for image classification," in *arXiv:1202.2745v1 [cs.CV]*, 2012.

[12] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A.C. Berg, and L. Fei-Fei, "Imagenet large scale visual recognition challenge," arXiv:1409.0575, 2014.

[13] F. Perronnin, Y. Liu, J. Sanchez, and H. Poirier, "Large-scale image retrieval with compressed fisher vectors," in *Computer Vision and Pattern Recognition*, June 2010, pp. 3384–3391.

[14] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," *CoRR*, 2014.

[15] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, 2014.

[16] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks," in *Computer Vision and Pattern Recognition*, 2014.

[17] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "Cnn features off-the-shelf: an astounding baseline for recognition," in *Computer Vision and Pattern Recognition*, 2014.

[18] Y. Gong, L. Wang, R. Guo, and S. Lazebnik, "Multi-scale orderless pooling of deep convolutional activation features," *arXiv:1403.1840 [cs.CV]*, 2014.

[19] O. Chum, A. Mikulik, M. Perdoch, and J. Matas, "Total recall ii: Query expansion revisited," in *Computer Vision and Pattern Recognition*, 2011.

[20] R. Arandjelović and A. Zisserman, "Three things everyone should know to improve object retrieval," in *International Conference on Computer Vision*, 2012.

[21] G. Tolias, Y. Avrithis, and H. Jegou, "To aggregate or not to aggregate: selective match kernels for image search," in *International Conference on Computer Vision*, 2013.

[22] H. Jégou and A. Zisserman, "Triangulation embedding and democratic aggregation for image search," in *Computer Vision and Pattern Recognition*, 2014.

[23] Y. Jia, "Caffe: An open source convolutional architecture for fast feature embedding," http://caffe.berkeleyvision.org/, 2013.

[24] A. Popescu, G. Grefenstette, and P.-A. Moellic, "Gazetiki: Automatic creation of a geographical gazetteer," in *Joint Conference on Digital libraries*, 2008.

[25] H. Jégou and O. Chum, "Negative evidences and cooccurrences in image retrieval: the benefit of pca and whitening," in *European Conference on Computer Vision*, 2012.